Co-designing an Energy Efficient System

Luigi Brochard
Distinguished Engineer , HPC&AI Lenovo
lbrochard@lenovo.com

MaX International Conference 2018
Trieste 29.01.2018

Lenovo™

Industry Thermal Challenges

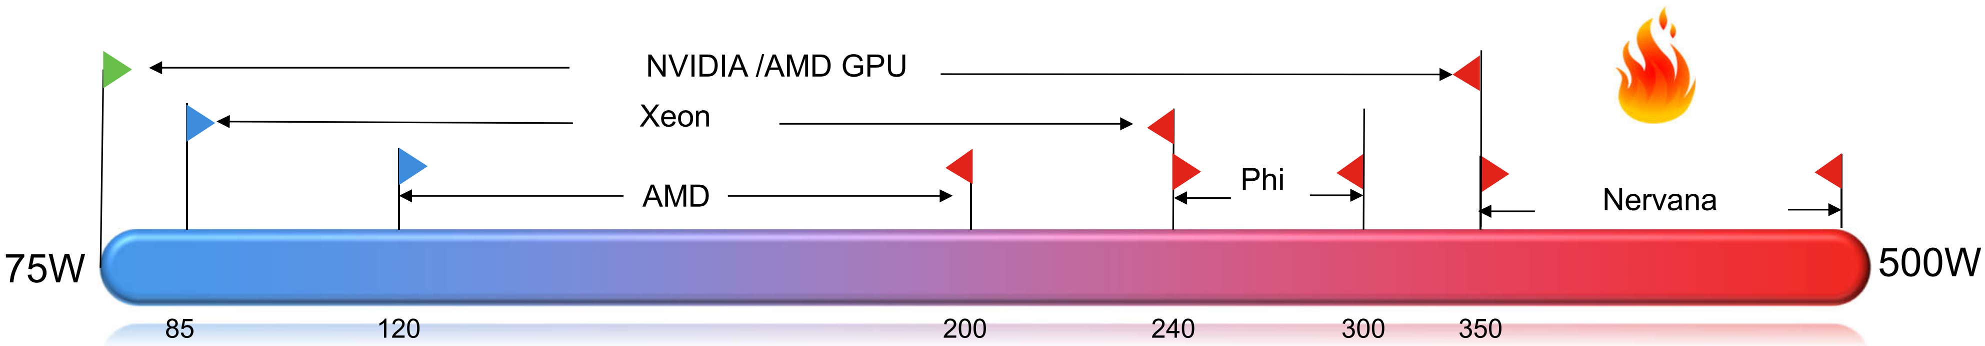

- Maintaining Moore's Law with increased competition is resulting in higher component power
- Increased memory count, NVMe adoption, and I/O requirements are driving packaging and feature tradeoffs
- To sustain increased performance servers will have to be less dense or use new cooling technology

How are we working on energy efficiency ?

- Higher Flops/Watt processor
- Water Cooling
- Software for power/energy management

Lenovo references with DWC (2012-2018)

Site	*Nodes*	*Country*	*Install date*	*Max inlet temperature*
LRZ SuperMUC	9298	Germany	2012	45°C
LRZ SuperMUC 2	3096	Germany	2014	45°C
LRZ SuperCool2	438	Germany	2015	50°C
NTU	40	Singapore	2012	45°C
Enercon	136	Germany	2013	45°C
US Army (Maui)	756	United States	2013	45°C
Exxon Research	504	United States	2014	45°C
NASA Goddard	80	United States	2014	45°C
PIK	312	Germany	2015	45°C
KIT	1152	Germany	2015	45°C
Birmingham U ph1	28	UK	2015	45°C
Birmingham U ph2	132	UK	2016	45°C
T-Systems	316	Germany	2016	45°C
MMD	296	Malaysia	2016	45°C
UNINET	964	Norway	2016	45°C
Peking U	204	China	2017	45°C
LPSC Trivandrum	72	India	2018	45°C
LRZ SuperMUC NG	6480	Germany	2018	50°C

More than 4.000 nodes with Lenovo DWC technology

2 X 3 petaflops SuperMUC systems at LRZ Phase 1 & Phase 2

Phase 1

Ranked 27 and 28 in TOP500 June 2016

- Fastest computer in Europe on TOP500, June 2012
  - 9324 nodes with 2 Intel Sandy Bridge EP CPUs
  - HPL = 2.9 petaflop/s
  - InfiniBand FDR10 interconnect
  - Large File Space for multiple purpose
    - 10 Petabyte File Space based on IBM GPFS with 200 GB/s I/O bandwidth
- Innovative technology for energy effective computing
  - Hot Water Cooling (45°C)
  - Energy Aware Scheduling
- Most energy efficient high-end HPC system
  - PUE 1.1
  - **Total power consumption over 5 years reduced by ~ 37% from 27.6 M€ to 17.4 M€**

Phase 2

- 3096 nx360 M5 compute nodes Haswell EP CPUs
- HPL = 2.8 petaflop/s
- Direct Hot Water Cooled, Energy Aware Scheduling
- InfiniBand FDR14
- GPFS, 10 x GSS26, 7.5 PB capacity , 100 GB/s I/O bandwidth

Three generation of direct water cooled systems

- iDataplex dx360M4 (2010-2013)
- NextScale nx360M5 WCT (2013-2016)
- OceanCat SD650 (2017 - 2018)

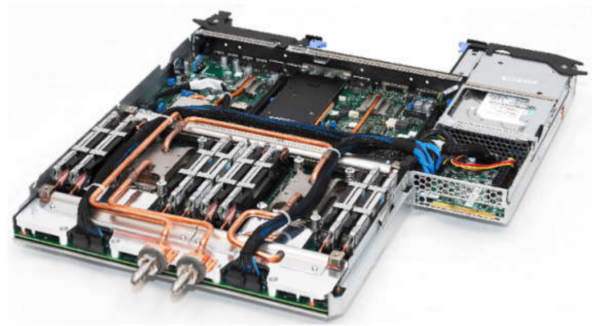

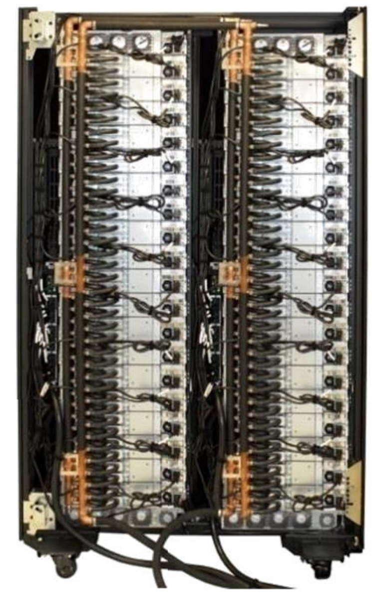

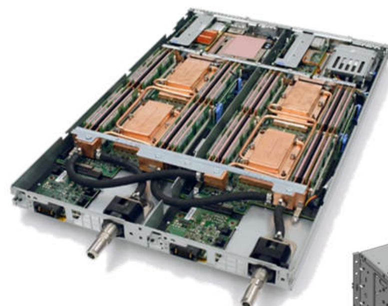

- Direct Water cooling CPU/DIMMS/VRs
  - **upto 90% of heat goes to water**
- Inlet water temperature
  - Up to 45-**50°C**

    => Free cooling all year long in most geo
- Wasted Heat Water is hot enough to be efficiently reused
  - like **with Adsorption chiller => ERE <<1**
- 3rd generation Water Cooling system in production
  - About 20.000 nodes installed

NextScale Chassis

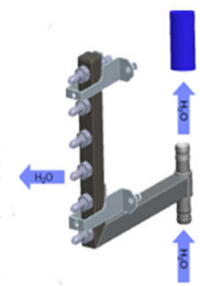

Scalable Manifold

PUE, ITUE and ERE

- PUE

$$PUE = \frac{\text{Total Facility Power}}{\text{IT Equipment Power}}$$

  - **Power usage effectiveness** (**PUE**) is a measure of how efficiently a computer data center uses its power;
  - PUE is the ratio of total power used by a computer facility[1] to the power delivered to computing equipment.
  - Ideal value is 1.0
  - It does not take into account how IT power can be optimised
- ITUE

$$ITUE = \frac{(\text{IT power} + \text{VR} + \text{PSU} + \text{Fan})}{\text{IT Power}}$$

  - **IT power effectiveness** (ITUE) measures how the node power can be optimised
  - Ideal value if 1.0
- ERE

$$ERE = \frac{\text{Total Facility Power} - \text{Treuse}}{\text{IT Equipment Power}}$$

  - **Energy Reuse Effectiveness** measures how efficient a data center reuses the power dissipated by the computer
  - ERE is the ratio of total amount of power used by a computer facility[1] to the power delivered to computing equipment.
  - An ideal ERE is 0.0. If no reuse, ERE = PUE

CooLMUC-2: Waste Heat Re-Use for Chilled Water Production

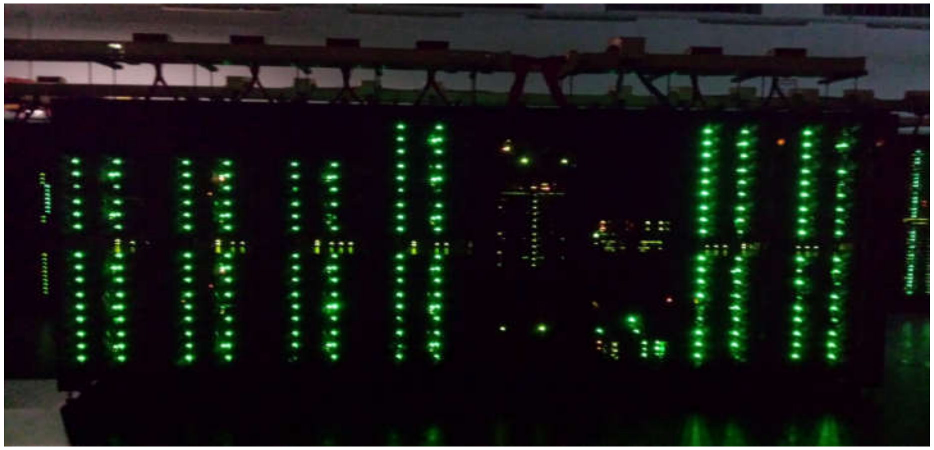

- **Lenovo NeXtScale Water Cool (WCT) system technology**
  - ✓ Water inlet temperatures 50 °C
  - ✓ All season chiller-less cooling
  - ✓ 384 compute nodes
  - ✓ 466 TFlop/s peak performance
- **SorTech Adsorbtion Chillers**
  - ✓ based of zeolite coated metal fiber heat exchangers
  - ✓ a factor 3 higher than current chillers based on silica gel
  - ✓ COP = 60%
  - ✓ ERE = 0.3

CooLMUC-2: ERE = 0.3

$$ERE = \frac{\text{Total Facility Power} - \text{Treuse}}{\text{IT Equipment Power}} = \frac{120 - 87}{104} = 0.32$$

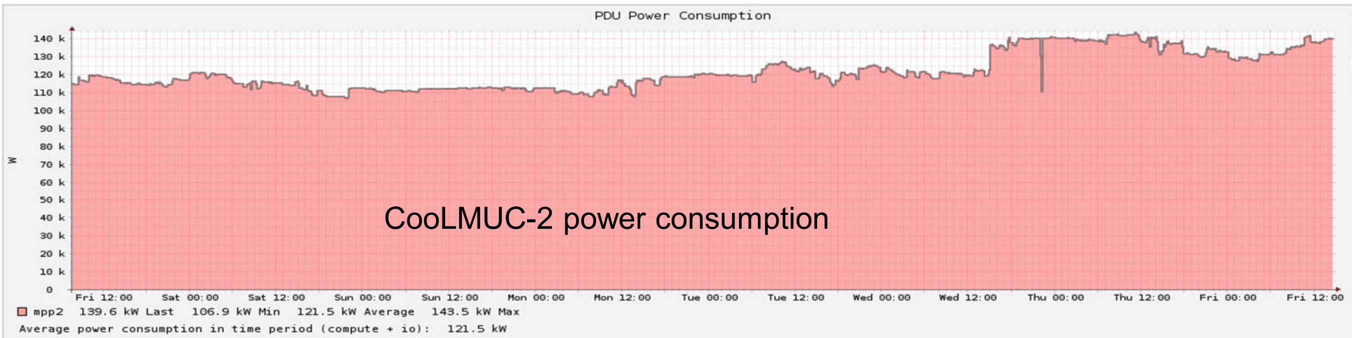

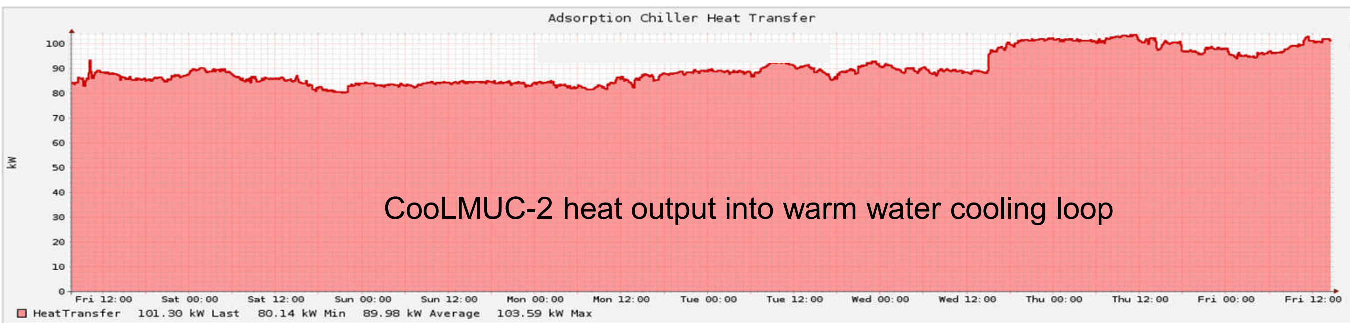

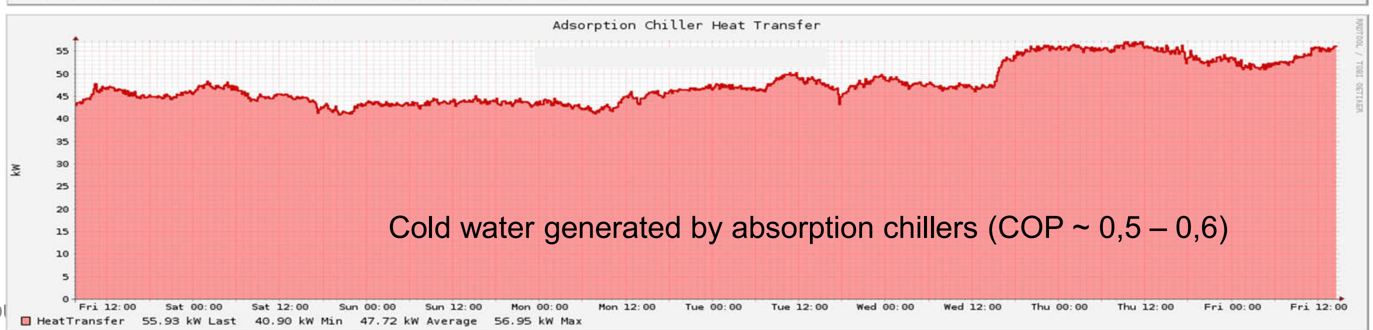

Example of energy cost with various cooling

- Air, RDHx and DWC energy cost
- **EnergyCost = *Total Power * Price per MW/year**
- **Total Power = (f_{air} * PUE_{air} + f_{cold} * PUE_{cold} + f_{warm} * PUE_{warm}) *Total IT Power**
  - Where f_{air}, f_{cold} and f_{warm} are the power consumption ratios using air, cold or hot water cooling vs total power
    - f_{air} = Total power cooled by air / Total Power
    - ..
- We assume:
  - 1MW server power consumption cooled with air or RDHX (cold water) and DWC with hot water at 50°C:
    - DWC: 90% heat to water leading to 10% heat to air residual or cold water
    - DWC: reduce power by 10% => 0.9 MW
      - Heat to water = 765 kW, Heat to air = 135 kW
  - 100 kW power consumption cooled with RDHx (storage, network…)
  - Total Power = 1.1 MW or 1 MW depending on cooling
  - PUE_{air} = 1.60, PUE_{cold} = 1.40 and PUE_{warm} = 1.06
  - Price of electricity is 1M€ (or $) per 1 MW per year

Lenovo

Example of energy cost with various cooling

- Partial heat reuse
  - Adsorption chillers can produce with COP = 0.6 chilled water with hot water and PUE of 1.06 to produce it
  - With DWC Hot water, Total $power_{warm}$ = 765 kW leading to 460 kW of chilled water capacity
  - Chilled water capacity is higher than chilled water needed
  - => partial heat reuse
- Full heat reuse
  - Data Center has more equipment which are cooled by CRAC, RDHX or In-Row coolers
  - Why not use the unused chilled water capacity for the Data Center others equipments ?

Energy Cost /year M € or $	no heat reuse	partial heat reuse	full heat reuse
Air only	1.76		
RDHx only	1.54		
DWC + air	1.19	1.13	0.91
DWC + RDHx	1.14	1.06	0.84

Relative to Air cooling	no heat reuse	partial heat reuse	full heat reuse
Air only	100%		
RDHx only	88%		
DWC + air	67%	64%	52%
DWC + RDHx	65%	60%	48%

Relative to DWC no reuse	no heat reuse	partial heat reuse	full heat reuse
Air only			
RDHx only			
DWC + air	100%	95%	77%
DWC + RDHx	96%	89%	70%

$$EnergyCost_{\text{with full heat reause}} \sim EnergyCost_{\text{with no heat reause}} * ERE / PUE$$

Lenovo

Value of Direct Water Cooling on SuperMUC

- Higher TDP processors (165 W)

- Reduced server power comsumption
  - Lower processor power consumption (~ 6%)
  - No fan per node (~ 4%)
- Reduce cooling power consumption
  - With DWC at 45°C, we assume free cooling all year long (~ 25%)

- Additional savings with xCAT and LL EAS

Total savings ~35-40%

SuperMUC NG system at LRZ

Phase 1

- Based on Xeon Skylake
  - 6334 Nodes with 2 Intel SKL @205 W CPUs
  - HPL ~ 20 PetaFLOP/s
  - OPA island based Interconnect
  - Large File Space on IBM Spectrum Scale
    - Scratch : 51 PB, 500GigaByte/s IOR bw
    - …
- Energy Effective Computing
  - More efficient Hot Water Cooling
  - Waste Heat Reuse
  - Dynamic Energy Aware Run time
- Best TCO and Energy Efficiency
  - PUE ~1.08
  - ERE ~0.30
  - **Total Power consumption over 5 years to be reduced upto ~50%**

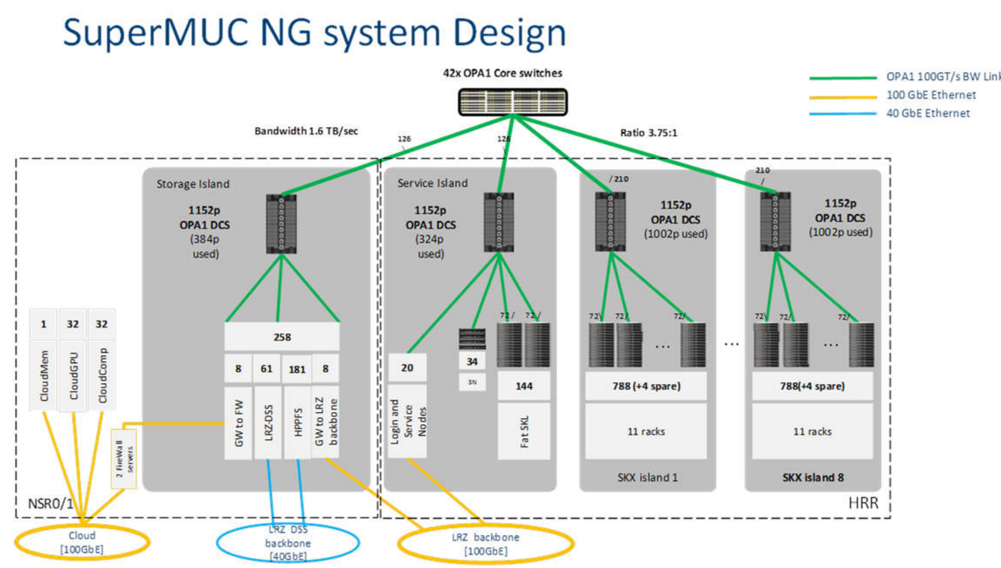

Value of Direct Water Cooling with SuperMUC NG

- Higher TDP processors (205+ W)
- Server power comsumption
  - Lower processor power consumption (~ 6%)
  - No fan per node (~ 4%)
- Cooling power consumption
  - With DWC upto 50°C, we assume free cooling all year long (~ 25%)

- Additional savings with xCAT and SLURM EAR

Savings = ~35-40%

- Heat Reuse
  - With DWC at 50°C, additional 30% savings as free chilled water is generated through adorption chillers

  With total heat reuse total savings => ~50%

How are we working on energy efficiency ?

- High Flops / Watt processor
- Water Cooling
- Software for energy management

How was this work done ?

LL-EAS phases to set optimal frequency for jobs

- **Initialization phase at LL installation time**
  - LL compute on all nodes the coefficents required for optimal frequency calculations
- **User submit a job**
  - User submit his job with a tag
  - Job is run at nominal frequency
  - In the background:
    - LL measures power, energy, time and hpm counters for the job
    - LL predicts power(i), energy(i), time (i) if job was run a different frequency i
  - LL writes Energy report for the job in the xCAT/LL DB
- **User resubmit a job with same tag**
  - Given the energy policy and the tag, LL determines optimal frequency j
  - LL set nodes for the job at frequency j
  - In the background:
    - LL measures power, energy, time and hpm counters for the job
    - LL compares measurement and prediction, and provide correction actions if needed
  - LL add new record with new energy report for the job in the xCAT/LL DB

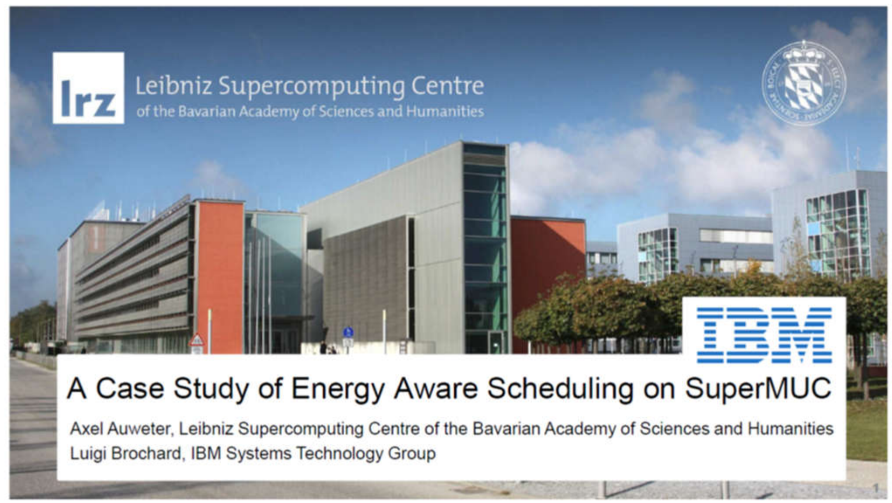

How was this work done ?

Energy Aware Run time

- Offer a dynamic and transparent solution to energy awareness :
  - Avoiding having to re-execute applications again and again
  - Easy to use
    - Without source code modifications
    - Without historic application information
    - Supporting standard programming models: MPI, MPI+OpenMP
    - Using standard libraries and tools as much as possible to be easily portable
    - Open Source
  - Frequency change based on simple Energy Policies with performance thresholds
  - Minimizing the overhead introduced

EAR high level view

- Automatic and dynamic frequency selection based on:
  - Distributed architecture / Low overhead
  - Architecture characterization (learning phase)
  - Application characterization
    - Outer loop detection (DPD)
    - Application signature computation (CPI,GBS,POWER,TIME)
  - Performance and power projection
  - Users/System policy definition for frequency selection (configured with thresholds)
    - MINIMIZE_ENERGY_TO_SOLUTION
      - Goal: To save energy by reducing frequency (with potential performance degradation)
      - We limit the performance degradation with a MAX_PERFORMANCE_DEGRADATION threshold
    - MINIMIZE_TIME_TO_SOLUTION
      - Goal: To reduce time by increasing frequency (with potential energy increase)
      - We use a MIN_PERFORMANCE_EFFICIENCY_GAIN threshold to avoid that application that do not scale with frequency to consume more energy for nothing
- GUI to monitor power and performance

LiCO – Lenovo intelligent Computing Orchestrator

A <u>single software stack</u> optimized for Intel and NVIDIA processors to efficiently manage both HPC & AI workloads

For HPC

- Easy-to-use interface for users to submit and manage jobs
- Full access to native tools in the stack for more technical users
- Built on an OpenHPC software base, with Lenovo value-add capabilities and optimizations

Simplify use and job management for HPC

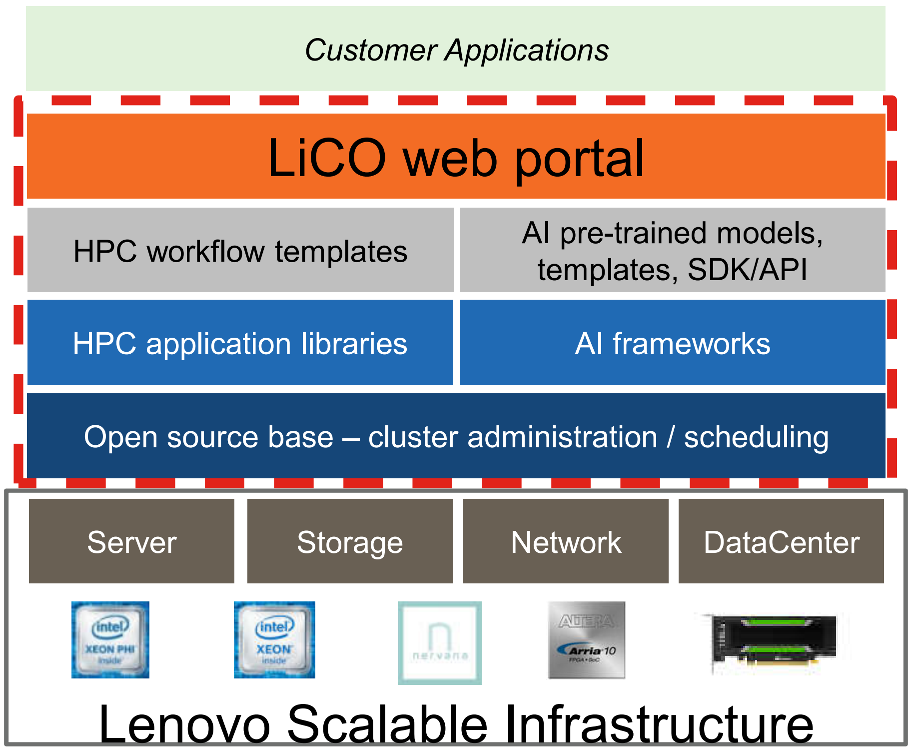

For AI

- Execute jobs, monitor training progress through a single GUI
- Easily try different frameworks, system types to determine best fit
- Out of the box scaling for both Intel and NVIDIA environments

Easy access to train and optimize AI models

Lenovo

LiCO for HPC

- Single GUI consolidates functionality for both administrators and users

- Built on an OpenHPC foundation
  - Monitoring, Scheduling, MPI, etc.

+

- Lenovo value-added capabilities, developed through client collaboration
  - Open web portal (Oxford & South Hampton)
  - Energy-Aware Runtime (BSC)

Validated stack of open tools to simplify cluster environments

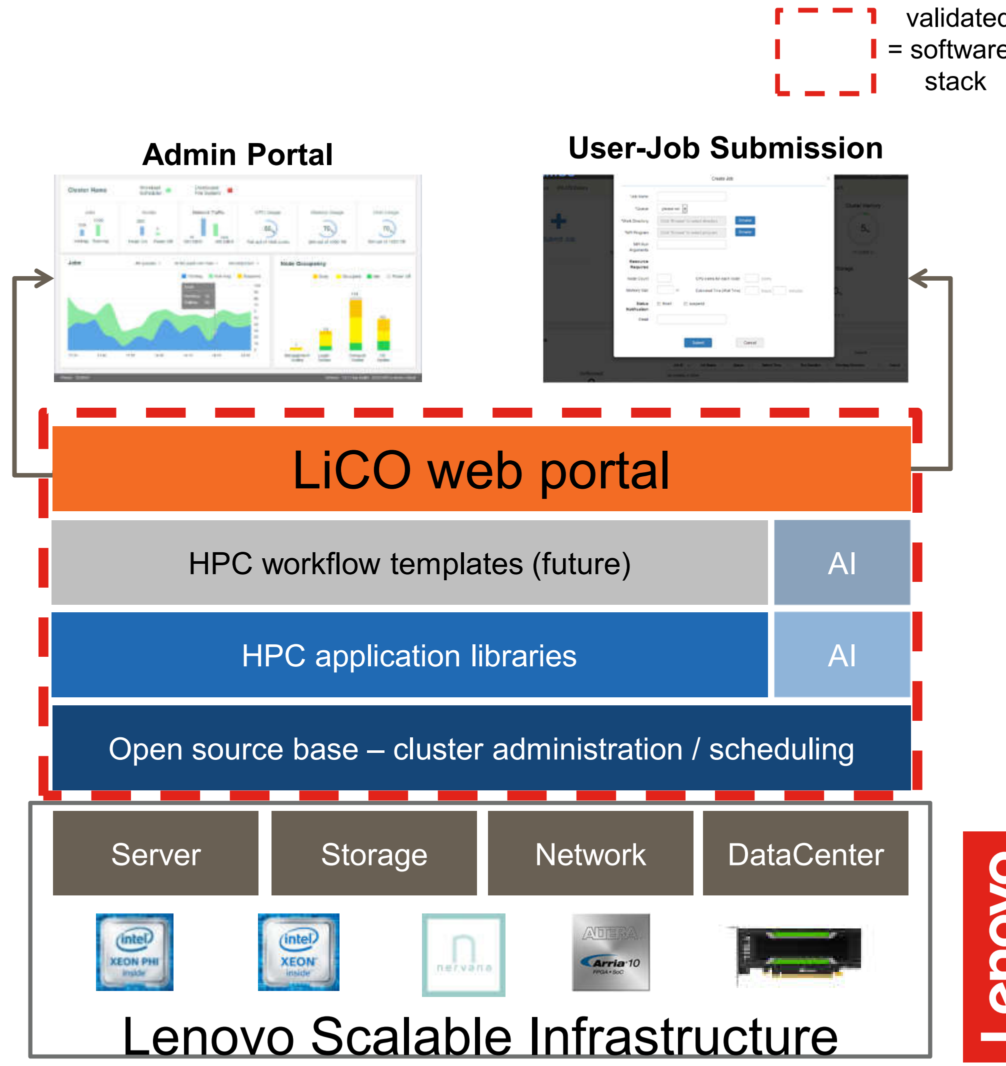

Conclusions

- Co-design with end users and research labs is crucial
- Example is Lenovo Energy Efficient systems
  - LRZ & BSC have been/are bringing their own skills to the development of new technologies

thanks.

Different is better